

Efficiently Using Second Order Information in Large ℓ_1 Regularization Problems

Xiaocheng Tang*

Katya Scheinberg†

March 28, 2013

Abstract

We propose a novel general algorithm LHAC that efficiently uses second-order information to train a class of large-scale ℓ_1 -regularized problems. Our method executes cheap iterations while achieving fast local convergence rate by exploiting the special structure of a low-rank matrix, constructed via quasi-Newton approximation of the Hessian of the smooth loss function. A greedy active-set strategy, based on the largest violations in the dual constraints, is employed to maintain a working set that iteratively estimates the complement of the optimal active set. This allows for smaller size of subproblems and eventually identifies the optimal active set. Empirical comparisons confirm that LHAC is highly competitive with several recently proposed state-of-the-art specialized solvers for sparse logistic regression and sparse inverse covariance matrix selection.

1 Introduction

We consider convex sparse unconstrained minimization problem of the following general form

$$\min_w F(w) = \lambda \|w\|_1 + L(w) \quad (1)$$

where $L : \mathbb{R}^p \rightarrow \mathbb{R}$ is convex and twice differentiable and $\lambda > 0$ is the regularization parameter that controls the sparsity of w . More generally, the regularization term $\lambda \|w\|_1$ can be extended to $\|\lambda \circ w\|_1 = \sum_{i=1}^p \lambda_i |w_i|$ to allow for different regularization weights on different entries, e.g., when there is a certain sparsity pattern desired in the solution w . We will focus on the simpler form as in (1) in this work for the sake of simplicity of presentation, as the extension to the general form is straightforward.

Problems of form (1) have been the focus of much research lately in the fields of signal processing and machine learning. This form encompasses a variety of machine learning models, in which feature selection is desirable, such as sparse logistic regression Yuan et al. [2010, 2011], Shalev-Shwartz and Tewari [2009], sparse inverse covariance selection Hsieh et al. [2011], Olsen et al. [2012], Scheinberg and Rish [2009], Lasso Tibshirani [1996], etc. These settings often present common difficulties to optimization algorithms due to their large scale. During the past decade most optimization effort aimed at these problems focused on development of efficient first-order methods, such as accelerated proximal gradients methods Nesterov [2007], Beck and Teboulle [2009], Wright et al. [2009], block

*Lehigh University, Bethlehem, PA 18015 USA. xct@lehigh.edu

†Lehigh University, Bethlehem, PA 18015 USA. katyas@lehigh.edu

coordinate descent methods Yuan et al. [2011], Friedman et al. [2010, 2008], Scheinberg and Rish [2009] and alternating directions methods Scheinberg et al. [2010]. These methods enjoy low per-iteration complexity, but typically have low local convergence rates. Their performance is often hampered by small step sizes. This, of course, has been known about first-order methods for a long time, however, due to the very large size of these problems, second order methods are often not a practical alternative. In particular, constructing and storing a Hessian matrix, let alone inverting it, is prohibitively expensive for values of p larger than 10000, which often makes the use of the Hessian in large-scale problems impractical, regardless of the benefits of fast local convergence rate.

Nevertheless, recently several new methods were proposed for sparse optimization which make careful use of second order information Hsieh et al. [2011], Yuan et al. [2011], Olsen et al. [2012], Byrd et al. [2012]. These methods explore the following special properties of the sparse problems: at the optimality many of the elements of w are expected to equal 0, hence methods which explore active set-like approaches can benefit from small sizes of subproblems. Whenever the subproblems are not small, these new methods exploit the idea that the subproblems do not need to be solved accurately. In particular we take a note of the following methods.

Yuan et al. [2011] proposes a specialized GLMNET Friedman et al. [2010] implementation for sparse logistic regression, where coordinate descent method is applied to the lasso subproblem constructed using the Hessian of $L(w)$ – the smooth component of the objective $F(w)$. Two major improvements are discussed to enhance GLMNET for larger problem – exploitation of the special structure of the Hessian to reduce the complexity of each coordinate step so that it is linear on the number of training instances, and a two-level shrinking scheme proposed to focus the minimization on smaller subproblems. Hsieh et al. [2011] later use the similar ideas in their specialized algorithm called QUIC for sparse inverse covariance selection. Benefiting from both its active-set strategy, which eventually converges to the optimal nonzero subspace, and its efficient use of the Hessian, QUIC behaves as Newton-like algorithms and is able to claim quadratic local convergence. Another related line of work begins with paper by Wright [2012], which proposes and analyzes an algorithm that is characterized by a two-phase minimization step for obtaining the improving direction, with one phase where a gradient descent step is taken towards minimizing the subproblem, and an enhanced phase where a Newton-like step is carried out in the nonzero subspace resulted from the first phase. Similar ideas are also explored in Olsen et al. [2012], in which a class of orthant-based Newton method is proposed such that Newton-like algorithm is applied in an orthant face which lies in a reduced subspace and is identified by first taking a steepest descent step. A backtracking line search, however, has to be put in place afterwards to project the step back onto the orthant whenever the step overshoots.

The above mentioned methods share similar attributes. In particular, all of them incorporate actual Hessian either by confining the subproblem minimization to a smooth subspace Wright [2012], Olsen et al. [2012], or by using it along with coordinate descent methods Yuan et al. [2011], Hsieh et al. [2011]. Active-set strategy is another key element shared by these approaches, which facilitates the use of the Hessian, often requiring only the *reduced Hessian* rather than the full one, and more importantly, help identify the optimal nonzero subspace and eventually achieve (with the Hessian) fast local convergence. Unfortunately, however, most of those active-set methods are only able to shrink the size of the subspace significantly when the current iterate is close enough to the optimality. Some algorithms are aware of this, e.g., Wright [2012] gives up the Hessian and returns to first-order steps if the size of the subspace exceeds 500, QUIC uses a small number ϵ (set

to a constant) to control the subspace size, etc. Hence, the efficient use of second order information in large problems is still a challenge. Of the aforementioned methods, the ones by Yuan et al. [2011] and Hsieh et al. [2011] produce the most satisfying results. But we note that both are specialized algorithms that heavily depends on the special structure present in the Hessians of the corresponding models. Hence their use will be limited.

In this work we make use of similar ideas as in Yuan et al. [2011] and Hsieh et al. [2011], but further improve upon these ideas to obtain efficient general schemes. In particular, we use a different working set selection strategy than those used in Hsieh et al. [2011]. We choose to select a working set by observing the largest violations in the dual constraints. Similar technique has been successfully used in optimization for many decades, for instance in Linear Programming Bixby et al. [1992] and in SVM Scheinberg [2006]. As in QUIC and GLMNET we optimize the Lasso subproblems using coordinate descent, but we estimate the Hessian using limited memory BFGS method Nocedal and Wright [2006] because the low rank form of the Hessian estimates reduce the per iteration complexity of each coordinate descent step to a constant.

Our goal here, as mentioned above, is to achieve second-order type convergence rate while maintaining a comparable per iteration complexity with that of most first-order methods. Following the approach as in Yuan et al. [2011] and Hsieh et al. [2011], we apply coordinate descent methods iteratively to the *lasso* subproblems constructed at the current point. The acceleration of subproblem minimization, therefore, depends on controlling either the number of coordinate descent steps or the complexity of each individual step. The contributions of our work are thus twofolds. First of all, we adaptively maintain a working set of coordinates with the largest optimality violations, such that the steps we take along those coordinates always provide the best objective function value improvement. The greedy nature of this approach helps reduce the violation of optimality conditions rather aggressively in practice while effectively avoiding zero updates, and more importantly, extends/shrinks (depending on the initial point) the working set incrementally until it converges to the complement of the optimal active set. Secondly, we explore the use of the Hessian and Hessian approximations in the coordinate descent framework. We show that each coordinate descent step can be greatly accelerated by the use of a special form of limited-memory BFGS method Nocedal and Wright [2006]. For example, in the case of sparse logistic regression, given the Hessian of logistic loss $L(w)$ as $B = CX^TDX$ where $D \in \mathbb{R}^{N \times N}$ is a diagonal matrix and $X \in \mathbb{R}^{N \times p}$ is the data matrix, the best implementation so far can only reduce the complexity of each coordinate descent step to $O(N)$ flops Yuan et al. [2011], while with the help of LBFGS which approximates the Hessian by a low rank matrix $B = \gamma I - Q\hat{Q}$ where $Q \in \mathbb{R}^{p \times 2m}$ and $\hat{Q} \in \mathbb{R}^{2m \times p}$, we are able to bring that complexity down to a constant time $O(m)$, depending on the limited-memory parameter m which is often chosen between 5 and 20. The key observation here is that the Hessian approximation obtained by LBFGS is low rank unlike the true Hessian, and that can be exploited to expedite the computation of $(Bd)_i$, the main expense of every coordinate descent step, by letting $(Bd)_i = \gamma d_i - q_i^T \hat{d}$, where q_i is the i -th row vector of the matrix Q and \hat{d} is cached and updated using one column of \hat{Q} for each step.

The paper is organized as follows. In Section 2 we explain how the subproblems are generated using LBFGS updates and working set selection. In Section 3 we explain how coordinate descent is applied to solve the subproblems with low rank Hessians. In Section 4 we present computations results on two instances of sparse logistic regression and five instances of sparse inverse covariance selection. The results demonstrate significant advantage of our approach compared to the other methods which inspired this work.

2 Outer iterations

Based on a generalization of the sequential quadratic programming method for nonlinear optimization Nocedal and Wright [2006], Tseng and Yun [2009], our approach iteratively constructs a piecewise quadratic model to be used in the step computation. At iteration k the model is obtained by expanding the smooth component $L(w)$ around the current iterate w_k and keeping the ℓ_1 regularization term, as follows:

$$d_k = \arg \min_{d, B_k \succ 0} \{ \nabla L_k^T d + d^T B_k d + \lambda \|w_k + d\|_1 \} \quad (2)$$

where B_k can be any positive definite matrix Tseng and Yun [2009]. In particular, we note that both Hsieh et al. [2011] and Yuan et al. [2011] choose B_k to be the Hessian of $L(w)$, in which case the objective function of (2) will be simply composed of the second-order Taylor expansion of $L(w)$ and the ℓ_1 term $\lambda \|w_k + d\|_1$.

An active-set method maintains a set of indices \mathcal{A} that iteratively estimates the optimal active set \mathcal{A}^* which contains indices of zero entries in the optimal solution w^* of (1)

$$\mathcal{A}^* = \{i \in \mathcal{P} \mid (w^*)_i = 0\} \quad (3)$$

where $\mathcal{P} = \{1, 2, \dots, p\}$. We use \mathcal{A}_k to denote the set \mathcal{A} at k -th iteration. For those coordinates $i \in \mathcal{A}_k$, we fix its corresponding entry in the current iteration $(w_k)_i$ such that it does not change from current iteration to the next $(w_k)_i = (w_{k+1})_i$. That is, equivalently, to say that the descent step along that coordinate stays zero $(d_k)_i = 0$. Adding the active-set strategy to our descent step computation (2), we thus obtain

$$\begin{aligned} d_k = \arg \min_d \{ & \nabla L_k^T d + d^T B_k d + \lambda \|w_k + d\|_1 \\ \text{s.t. } & B_k \succ 0, d_i = 0, \forall i \in \mathcal{A}_k \} \end{aligned} \quad (4)$$

Next we are going to discuss the particular choice we make in selecting the positive definite matrix B_k and \mathcal{A}_k , or its complement $\mathcal{I}_k = \{i \in \mathcal{P} \mid i \notin \mathcal{A}_k\}$, which we refer to in this paper as the working set.

2.1 Low-Rank Hessian Approximation B_k

We make use of Theorem 1 from Nocedal and Wright [2006], which gives a specific form of the low-rank Hessian estimates, which we denote by B_k . B_k is essentially a low-rank approximation of the Hessian of $L(w)$ through the well-known limited-memory BFGS method, which allows the capture of the curvature information to help achieve a faster local convergence.

Theorem 1. *Let B_0 be symmetric and positive definite, and assume that the k vector pairs $\{s_i, t_i\}_{i=0}^{k-1}$ satisfy $s_i^T t_i > 0$, $s_i = w_{i+1} - w_i$ and $t_i = \nabla L_{i+1} - \nabla L_i$. Let B_k be obtained by applying k BFGS updates with these vector pairs to B_0 , using the formula:*

$$B_k = B_{k-1} - \frac{B_{k-1} s_{k-1} s_{k-1}^T B_{k-1}}{s_{k-1}^T B_{k-1} s_{k-1}} + \frac{y_{k-1} y_{k-1}^T}{y_{k-1}^T s_{k-1}} \quad (5)$$

We then have that

$$B_k = B_0 - [B_0 S_k \quad T_k] \begin{bmatrix} S_k^T B_0 S_k & L_k \\ L_k^T & -D_k \end{bmatrix}^{-1} \begin{bmatrix} S_k^T B_0 \\ T_k^T \end{bmatrix}, \quad (6)$$

where S_k and T_k are the $p \times k$ matrices defined by

$$S_k = [s_0, \dots, s_{k-1}], \quad T_k = [t_0, \dots, t_{k-1}], \quad (7)$$

while L_k and D_k are the $k \times k$ matrices

$$(L_k)_{i,j} = \begin{cases} s_{i-1}^T t_{j-1} & \text{if } i > j \\ 0 & \text{otherwise,} \end{cases} \quad (8)$$

$$D_k = \text{diag}[s_0^T t_0, \dots, s_{k-1}^T t_{k-1}]. \quad (9)$$

For large-scale problems, BFGS method is often used in the limited-memory setting, known as the L-BFGS method Nocedal and Wright [2006]. Note that the matrices S_k and T_k are augmented by one column every iteration according to (7), and the update (6) will soon become inefficient if all the previous pairs $\{s_i, t_i\}$ are used.

In the limited-memory setting, we maintain the set $\{s_i, t_i\}$ with the m most recent correction pairs by removing the oldest pair and adding a newly generated pair. Hence when the number of iteration k exceeds m , the representation of the matrices S_k, T_k, L_k, D_k needs to be slightly modified to reflect the changing nature of $\{s_i, t_i\}$, and S_k, T_k are maintained as the $p \times m$ matrices. Also, L_k and D_k are now the $m \times m$ matrices.

In the L-BFGS algorithm, the basic matrix B_0 may vary from iteration to iteration. A popular choice in practice is to set the basic matrix at k -th iteration to $B_0^{(k)} = \gamma_k I$, where

$$\gamma_k = \frac{t_{k-1}^T t_{k-1}}{t_{k-1}^T s_{k-1}} \quad (10)$$

which is proved effective in ensuring that the search direction is well-scaled so that less time is spent on line search. With this particular choice of $B_0^{(k)}$, we define Q and R to be

$$Q = [\gamma_k S_k \quad T_k] \quad (11)$$

$$R = \begin{bmatrix} \gamma_k S_k^T S_k & L_k \\ L_k^T & -D_k \end{bmatrix}^{-1} \quad (12)$$

and the formula to update B_k becomes

$$B_k = \gamma_k I - Q R Q^T \quad (13)$$

$$= \gamma_k I - Q \hat{Q} \quad \text{with } \hat{Q} = R Q^T \quad (14)$$

where γ_k is given by (10). Hence, the work to compute B_k only requires the matrix Q which is a collection of previous iterates and gradient differences, and R^{-1} that is a $2m$ by $2m$ nonsingular (as long as $s_i^T t_i > 0 \forall i$) matrix whose inverse is thus easy to compute given m is a small constant.

2.2 Greedy Active-set Selection $\mathcal{A}_k(\mathcal{I}_k)$

An obvious choice of \mathcal{I}_k would be $\mathcal{I}_k = \{1, \dots, p\}$, taking into account all the coordinates in subproblem minimization. But as we said earlier, this can be inefficient because there will potentially be many non-progressive steps where z in (24) ends up being zero. For example, if $d = 0$, $(B_k d)_j$ turns into zero and (24) becomes

$$z = \begin{cases} \frac{(\nabla L_k)_j + \lambda}{-(B_k)_{jj}} & \text{if } (\nabla L_k)_j + \lambda \leq (B_k)_{jj}(w_k)_j, \\ \frac{(\nabla L_k)_j - \lambda}{-(B_k)_{jj}} & \text{if } (\nabla L_k)_j - \lambda \geq (B_k)_{jj}(w_k)_j, \\ -(w_k)_j & \text{otherwise.} \end{cases} \quad (15)$$

which is equal to zero if the j th entry of the subgradient ∂F , defined below, is zero

$$(\partial F(w_k))_j = \begin{cases} (\nabla L(w_k))_j + \lambda & (w_k)_j > 0, \\ (\nabla L(w_k))_j - \lambda & (w_k)_j < 0, \\ \max\{|\nabla L(w_k))_j| - \lambda, 0\} & (w_k)_j = 0 \end{cases} \quad (16)$$

Hence, the possibly worst choice of \mathcal{I}_k would be letting $\mathcal{I}_k = \{i \in \mathcal{P} \mid (\partial F(w_k))_i = 0\}$, which will result in nothing but all zeros in coordinate descent update.

Let us now consider two potentially good choice of \mathcal{I}_k

$$\mathcal{I}_k^{(1)} = \{i \in \mathcal{P} \mid (\partial F(w_k))_i \neq 0\} \quad (17)$$

used by Yuan et al. [2011] and Hsieh et al. [2011] and

$$\mathcal{I}_k^{(2)} = \{i \in \mathcal{P} \mid (w_k)_i \neq 0\} \quad (18)$$

considered by Wright [2012] and Olsen et al. [2012]. Note that $\mathcal{I}_k^{(1)} \supseteq \mathcal{I}_k^{(2)}$, because in practice $(\partial F(w_k))_i$ can only be exactly zero if $(w_k)_i = 0$ according to (16), so we have that if $(w_k)_i \neq 0$ then $(\partial F(w_k))_i \neq 0$. Also note that the two sets will both converge to the optimal active set $\mathcal{I}_k^{(1)} = \mathcal{I}_k^{(2)} = \mathcal{I}^*$ when the correct non-zeros have been identified in w_k and w_k is close enough to the optimal w^* , because in which case the violations in the dual constraints of those zero entries in w_k will be zero such that $\mathcal{A}_k^{(1)} = \{i \in \mathcal{P} \mid (\partial F(w_k))_i = 0\} = \mathcal{A}_k^{(2)} = \{i \in \mathcal{P} \mid (w_k)_i = 0\}$. The above two observations can also be understood by representing $\mathcal{I}_k^{(1)}$ as the union of two sets $\mathcal{I}_k^{(1a)}$ and $\mathcal{I}_k^{(1b)}$, defined by

$$\mathcal{I}_k^{(1a)} = \{i \in \mathcal{P} \mid (\partial F(w_k))_i \neq 0 \text{ and } (w_k)_i \neq 0\} \quad (19)$$

$$\mathcal{I}_k^{(1b)} = \{i \in \mathcal{P} \mid (\partial F(w_k))_i \neq 0 \text{ and } (w_k)_i = 0\} \quad (20)$$

and we have that $\mathcal{I}_k^{(2)} = \mathcal{I}_k^{(1a)}$ in general, and that $\mathcal{I}_k^{(2)} = \mathcal{I}_k^{(1)}$ only when $\mathcal{I}_k^{(1b)} = \emptyset$.

Now let us introduce our rule to select the working set \mathcal{I}_k . Particularly, Let us use $\hat{i} := \mathcal{I}_k^{(1)}(i)$ to denote the coordinate $\hat{i} \in \mathcal{I}_k^{(1)}$ that has the i th largest violations $|(\partial F_k)_i|$ in the dual constraints. We then choose the working set for k -th iteration by

$$\mathcal{I}_k = \mathcal{I}_k^{(2)} \cup \bigcup_{i=1}^{s_k} \hat{i} \quad (21)$$

where s_k is a small integer chosen as a fraction of $|\mathcal{I}_k^{(1)}|$. Note that our \mathcal{I}_k is largely based on $\mathcal{I}_k^{(2)}$, whose size, as discussed above, is smaller than $\mathcal{I}_k^{(1)}$ as used by Yuan et al. [2011] and Hsieh et al. [2011]. This, of course, enables us to solve a smaller subproblem (4). However, we also note that choosing $\mathcal{I}_k = \mathcal{I}_k^{(2)}$ in our case is a bad idea because it does not allow zero elements of w to become nonzero, so the method may fail to converge. To ensure convergence, we have to let every coordinate enter or leave our working set \mathcal{I}_k , which is one purpose of the union of the set $\bigcup_{i=1}^{s_k} \hat{i}$.

2.3 Line Search and Convergence Analysis

After the step d_k is computed, a line search procedure needs to be employed in order for the convergence analysis to follow from the framework by Tseng and Yun [2009]. In this

work, we adapt the Armijo rule, choosing the step size α_k to be the largest element from $\{\beta^0, \beta^1, \beta^2, \dots\}$ satisfying

$$F(w_k + \alpha_k d_k) \leq F(w_k) + \alpha_k \sigma \Delta_k \quad (22)$$

where $0 < \beta < 1, 0 < \sigma < 1$, and $\Delta_k := \nabla L_k^T d_k + \lambda \|w_k + d_k\|_1 - \lambda \|w_k\|_1$. The convergence analysis from Tseng and Yun [2009] also requires the positive definiteness of the matrix B_k . Since we obtain B_k using LBFGS update, it is guaranteed to be positive definite as long as the product $s_i^T t_i > 0$. We note that when $L(w)$ in (1) is strongly convex, then $s_i^T t_i > 0$ holds at any two points, otherwise if $L(w)$ might be non-convex, then damped BFGS updating needs to be employed Nocedal and Wright [2006].

Algorithm 2.1 LHAC: LOW RANK HESSIAN APPROXIMATION IN ACTIVE-SET COORDINATE DESCENT

- 1: Choose an initial iterate w_0 , the LBFGS parameter m , the working set parameter s
 - 2: Set S, T, D, L to empty matrix
 - 3: Set $w_1 \leftarrow w_0, \gamma_1 \leftarrow 1, Q \leftarrow 0, \hat{Q} \leftarrow 0$
 - 4: Set the iteration counter $k \leftarrow 1$
 - 5: Compute the gradients ∇L_k and ∂F_k
 - 6: **while** optimality test returns false **do**
 - 7: Set $d_k \leftarrow 0$ and $\hat{d} \leftarrow 0$
 - 8: Compute the *working* set $\mathcal{I}_k = \mathcal{I}_k^{(2)} \cup \bigcup_{i=1}^s \hat{i}$
 - 9: **for** each coordinate $j \in \mathcal{I}_k$ **do**
 - 10: Compute $(B_k)_{jj} = \gamma_k - q_j^T \hat{Q}_j$
 - 11: **end for**
 - 12: **for** each coordinate $j \in \mathcal{I}_k$ **do**
 - 13: Compute $(B_k d)_j = \gamma_k d_j - q_j^T \hat{d}$
 - 14: Compute z according to (24)
 - 15: Update $d_k \leftarrow d_k + z e_j, \hat{d} \leftarrow \hat{d} + z \hat{Q}_j$
 - 16: **end for**
 - 17: Compute the step size α_k according to (22)
 - 18: Set $w_{k+1} \leftarrow w_k + \alpha_k d_k$
 - 19: Compute the gradients ∇L_{k+1} and ∂F_{k+1}
 - 20: Set $t_k \leftarrow \nabla L_{k+1} - \nabla L_k$ and $s_k \leftarrow w_{k+1} - w_k$
 - 21: Update S, T, D, L according to Theorem 1
 - 22: Compute Q, \hat{Q} according to Theorem 1
 - 23: **end while**
-

3 The inner problem

At k -th iteration given the current iterate w_k , we apply coordinate descent method to the piecewise quadratic subproblem (4) to obtain the direction d_k . Suppose j th coordinate in d is updated, hence $d' = d + z e_j$ (e_j is the j -th vector of the identity). Then z is obtained by solving the following one-dimensional problem

$$\min_z \frac{1}{2} (B_k)_{jj} z^2 + ((\nabla L_k)_j + (B_k d)_j) z + \lambda |(w_k)_j + (d)_j + z| \quad (23)$$

which is well known to have the following closed-form solution that can be obtained by one soft-thresholding step Donoho [1995], Hsieh et al. [2011]:

$$z = -c + \mathcal{S}(c - b/a, \lambda/a) \quad (24)$$

with a, b, c chosen to be $a = (B_k)_{jj}$, $b = (\nabla L_k)_j + (B_k d)_j$ and $c = (w_k)_j + (d)_j$.

As mentioned above, the special low-rank update of B_k provides us an opportunity to accelerate the coordinate descent process. To see how, let us recall that $B_k = \gamma_k I - Q\hat{Q}$, where $Q \in \mathbb{R}^{p \times 2m}$, $\hat{Q} \in \mathbb{R}^{2m \times p}$ and m chosen as a small constant. Clearly we do not need to explicitly store or compute B_k . Instead, since B_k is only used through $(B_k)_{ii}$ and $(B_k d)_i$ when applying soft-thresholding steps to updating each coordinate i , we can only store the diagonal elements of B_k and compute $(B_k d)_i$ on the fly whenever it is needed. Specifically,

$$(B_k)_{ii} = \gamma_k - q_i^T \hat{q}_i \quad (25)$$

where q_i is the i th row of the matrix Q and \hat{q}_i is the i th column vector of the matrix \hat{Q} . To compute $(B_k d)_i$, we maintain a $2m$ dimensional vector $\hat{d} := \hat{Q}d$, then

$$(B_k d)_i = \gamma_k d_i - q_i^T \hat{d} \quad (26)$$

which takes $O(2m)$ flops, instead of $O(p)$ if we multiply d by the i th row of B_k . Notice though, that after taking the soft-thresholding step \hat{d} has to be updated each time by $\hat{d} \leftarrow \hat{d} + z_i \hat{q}_i$. This update requires little effort given that \hat{q}_i is a vector with $2m$ dimensions where m is often chosen to be between 3 and 20. However, we need to use extra memory for caching \hat{Q} and \hat{d} which takes $O(2mp + 2m)$ space. With the other $O(2p + 2mp)$ space for storing the diagonal of B_k , Q and d , altogether we need $O(4mp + 2p + 2m)$ space, which can be written as $O(4mp)$ when $p \gg m$.

4 Experiments

4.1 Sparse Logistic Regression

The objective function of sparse logistic regression is given by

$$F(w) = \lambda \|w\|_1 + \frac{1}{N} \sum_{n=1}^N \log(1 + \exp(-y_n \cdot w^T x_n))$$

where $L(w) = \frac{1}{N} \sum_{n=1}^N \log(1 + \exp(-y_n \cdot w^T x_n))$ is the average logistic loss function and $\{(x_n, y_n)\}_{n=1}^N \in (\mathbb{R}^p \times \{-1, 1\})$ is the training set. The number of instances in the training set and the number of features are denoted by N and p respectively. Note that the evaluation of F requires $O(pN)$ flops and to compute the Hessian requires $O(Np^2)$ flops. Hence, we choose such training sets for our experiment that N and p are large enough to test the scalability of the algorithms and yet small enough to be able to run on a workstation.

We test the algorithms on both sparse and dense data sets. The statistics of those data sets used in the experiments are summarized in Table 1. In particular, one data set, RCV1 Lewis et al. [2004], is a text categorization test collection made available from Reuters Corpus Volume I (RCV1), an archive of over 800,000 manually categorized newswire stories. Each RCV1 document used in our experiment has been tokenized, stopworded, stemmed and represented by 47,236 features in a final form as weighted vectors. Another data set is

GISETTE, originally a handwriting digit recognition problem from NIPS 2003 feature selection challenge. Its feature set of size 5000 has been constructed in order to discriminate between two confusable handwritten digits: the four and the nine.

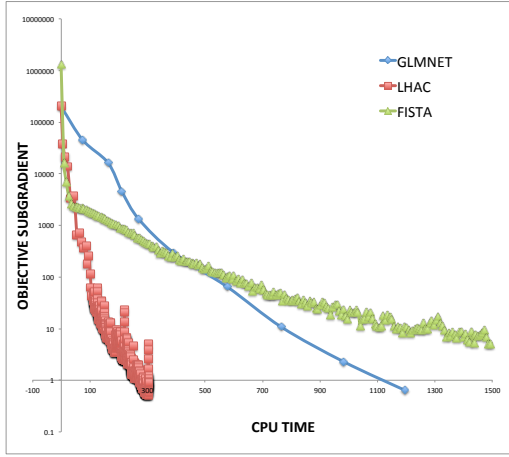
We compare LHAC with the first order method FISTA Beck and Teboulle [2009], and GLMNET Friedman et al. [2010], Yuan et al. [2011] which uses Hessian information. GLMNET in particular, originally proposed by Friedman et al. [2010] and later improved specifically for sparse logistic regression by Yuan et al. [2011] has been shown by Yuan et al. [2010] and Yuan et al. [2011] to be the state-of-the-art for sparse logistic regression (their experiments include algorithm such as OWL-QN Andrew and Gao [2007], TRON Lin and Mor [1999], SCD Shalev-Shwartz and Tewari [2009], BBR Genkin et al. [2007], etc.), so we compare LHAC with GLMNET. We implemented all three algorithms, for fair comparison, in MATLAB, and all the experiments were executed on AMD Opteron 2.0 GHz machines with 32G RAM and Linux OS.

We have generated eight training sets, with different number of training instances, from the two data sets RCV1 and GISETTE, four for each one. For RCV1, the training size increases from 500 to 2500; for GISETTE, it ranges from 500 to 5000. In all the experiments, we terminate the algorithm when the current objective subgradient is detected to be less than the precision ϵ times the initial subgradient. For each training set, we solve the problem twice using each algorithm, one with a large epsilon, e.g., $\epsilon = 10^{-2}$, to test an algorithm’s ability to quickly obtain a useful solution, the other one with a small epsilon, e.g., $\epsilon = 10^{-5}$, to see whether an algorithm can achieve an accurate solution. We report the runtime of each algorithm for all 16 experiments in Table 1. We also illustrate the results in Figure 1.

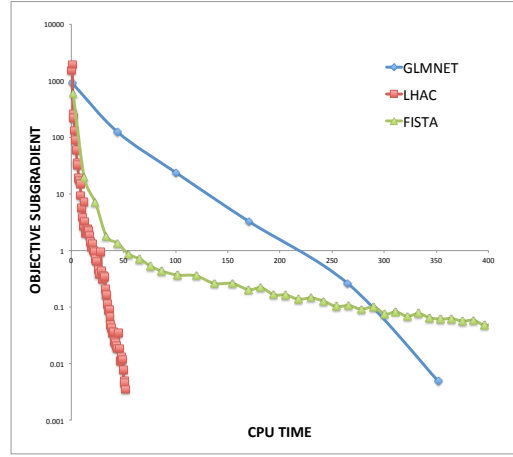
As can be seen from the results, the advantage of LHAC really lies in the fact that it absorbs the benefits from both the first order and the second order methods, such that it iterates fast, with low per iteration cost, and that it converges fast like other quasi-Newton methods, by taking advantage of the objective curvature information. For example, let us look at how FISTA compares with GLMNET in Table 1, which provides a good case to study the well-known trade-offs between first order and second order methods. Particularly, it can be seen that in most cases when ϵ is set to 10^{-2} , FISTA takes up much less time to terminate than GLMNET. However, FISTA falls short when high precision is demanded, e.g., ϵ set to 10^{-5} , in which case GLMNET almost always terminates faster than FISTA. Notably, LHAC is able to perform well in both cases. In fact, it is overwhelmingly faster than the other two in all but one experiment (the runtime difference between LHAC and GLMNET is close when the training size is sufficiently small), which brings us to another interesting aspect about LHAC. That is, the larger the size of the training set is, the larger the margin becomes between LHAC and GLMNET. We illustrate this observation in Figure 1(c) and 1(d), where the runtime of each algorithm ($\epsilon = 10^{-5}$) is plotted against the training size. Note that the complexity of LHAC scales almost linearly with the problem size (with a much flatter rate than FISTA), while that of GLMNET increases nonlinearly. Figure 1(a) and 1(b) plot the change of objective subgradient against elapsed cputime in one experiment. Again, it can be observed that LHAC iterates as efficient as FISTA in the beginning, and while FISTA gradually slows down near the optimality, it continues to work as GLMNET until reaching the optimality tolerance 10^{-5} .

4.2 Sparse Inverse Covariance Matrix Estimation

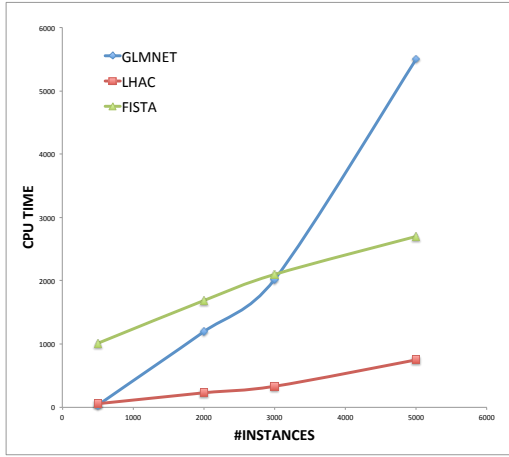
In this section we compare our algorithm LHAC with QUIC by Hsieh et al. [2011], a specialized solver that has been shown by Olsen et al. [2012] and Hsieh et al. [2011] to be



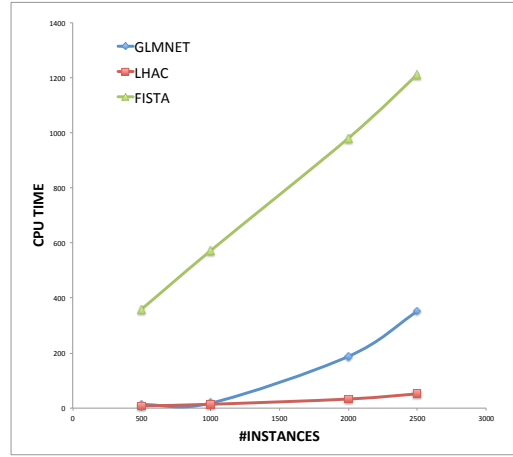
(a) GISETTE dense. (#Instances = 2000).



(b) RCV1 sparse. (#Instances = 2500).



(c) GISETTE dense. #Instances from 500 to 5000.



(d) RCV1 sparse. #Instances from 500 to 2500.

Figure 1: Convergence plots in 1(a) and 1(b) (the y-axes on log scale). Scalability plots in 1(c) and 1(d).

the state of the art for sparse covariance selection problem defined below

$$\min_{X \succ 0} f(X) = -\log \det X + \text{tr}(SX) + \lambda \|X\|_1 \quad (27)$$

where the optimization variables are in a matrix form $X \in \mathbb{R}^{p \times p}$ that is required to be positive definite.

For this experiment, we are interested in comparing the total complexity of the two algorithms. CPU time, as often used, will not be a reasonable complexity measure in this case, because QUIC is implemented in C++ and LHAC - in MATLAB. Instead, we decide to count the number of flops required by each algorithm to solve problem (27). Here we note that the work required by both algorithm consists mainly two parts – one for solving the subproblems and the other for computing Cholesky factorizations of X during line search, and that the time either algorithm spends on the first part – solving the subproblems – generally takes around 95 ~ 99% of the total elapsed time, as observed in the experiments. Hence we focus our comparisons on the first part of the work – the one required by applying coordinate descent to solving subproblems, and we note that

Table 1: CPU time comparisons on real world classification data sets. Data set *rcv1* has 47,326 features and 0.17% non-zeros; Data set *gisette* has 5,000 features and 99% non-zeros. *rcv1*₅₀₀ denotes that the training set is *rcv1* and 500 training samples are used. ϵ indicates the optimization tolerance on the objective subgradient. The results show that LHAC is significantly faster than other methods.

DATA SET	ϵ	CPUTIME(IN SECONDS)		
		FISTA	GLMNET	LHAC
<i>rcv1</i> ₅₀₀	10^{-2}	4.54	6.39	2.68
	10^{-5}	358.72	12.04	8.09
<i>rcv1</i> ₁₀₀₀	10^{-2}	14.28	6.18	3.53
	10^{-5}	572.41	18.13	14.13
<i>rcv1</i> ₂₀₀₀	10^{-2}	15.28	97.87	6.13
	10^{-5}	980.72	188.48	33.05
<i>rcv1</i> ₂₅₀₀	10^{-2}	17.71	170.13	9.95
	10^{-5}	1212.92	351.60	52.00
<i>gisette</i> ₅₀₀	10^{-2}	52.91	10.93	3.23
	10^{-5}	1009.66	35.95	54.90
<i>gisette</i> ₂₀₀₀	10^{-2}	85.41	269.26	35.02
	10^{-5}	1686.69	1195.48	229.67
<i>gisette</i> ₃₀₀₀	10^{-2}	101.52	507.65	40.17
	10^{-5}	2241.95	2021.40	332.83
<i>gisette</i> ₅₀₀₀	10^{-2}	102.78	1758.98	112.11
	10^{-5}	2661.95	5935.72	752.87

although LHAC generally requires more work in line search than QUIC (due to more outer iterations), it adds little to the total work for reasons stated above.

Now let us describe how we compute the total complexity. Let us use κ_o to denote the number of outer iteration and κ_i to denote the number of inner coordinate sweep. Let T_{CD} stand for the number of flops one coordinate descent step takes. Then we define the total complexity by

$$Complexity = \kappa_o \kappa_i p T_{CD} \quad (28)$$

In theory, the two algorithms have similar κ_i because they both apply coordinate descent to a lasso subproblem, but different κ_o and T_{CD} . Particularly, κ_o of QUIC is smaller than that of LHAC whereas T_{CD} of LHAC is smaller than that of QUIC. The reason is that QUIC uses the actual Hessian matrix of the smooth part $\text{tr}(SX) - \log \det X$ rather than the low-rank Hessian approximation as in LHAC, which results in a different convergence rate (κ_o) and also a different complexity (T_{CD}) in computing the coordinate descent step (24). As we discussed earlier, the special structure of the low-rank matrix enables LHAC to accelerate that step to $O(m)$ flops with m a constant number (chosen as 10 for this experiment). Whereas in the case of QUIC, one coordinate step takes problem-dependent complexity $O(p)$ where p is the dimension of X , but it achieves quadratic convergence when close to the optimality. The above observations make it interesting to see in practice how the *Complexity*, defined in (28), will compare for the two algorithms since it is not obvious which one is better through theoretical analysis alone.

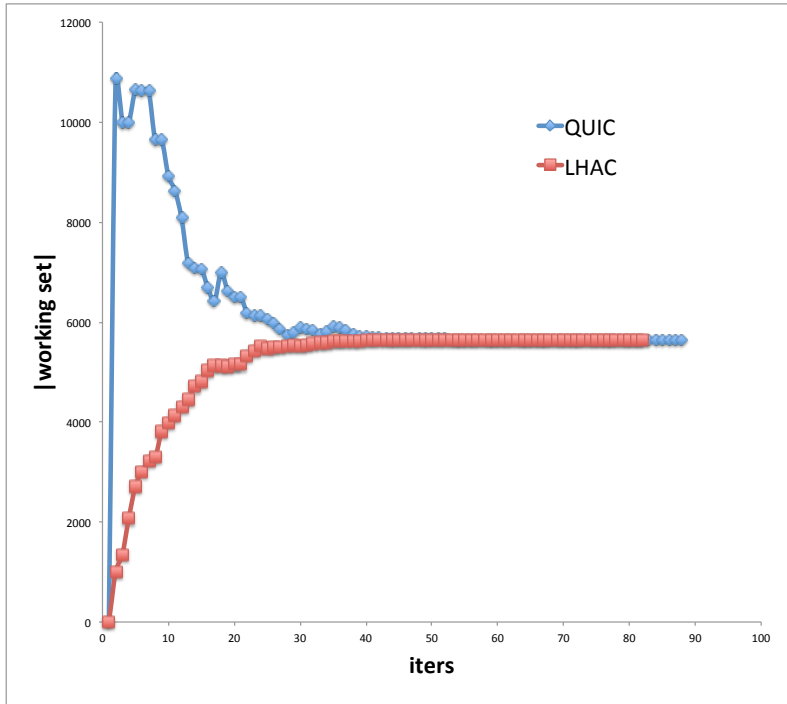


Figure 2: Comparing two strategies - the strategy used by QUIC and the one used by LHAC - of choosing the working set on ER data set. Note that the working set maintained by LHAC approaches the optimal size from below while that of QUIC from above.

The results are presented in Table 2, and in Figure 2 where we compare our active-set strategy with the one used by QUIC. The way we compute the *Complexity* is by counting the number of flops in the coordinate descent step in each algorithm. For QUIC, we use their C++ implementation Hsieh et al. [2011] and add a counter directly in their code; for LHAC, we use our own MATLAB implementation. We report the results on real world data from gene expression networks preprocessed by Li and Toh [2010]. We set the regularization parameter $\lambda = 0.5$ for all the experiments as suggested in Li and Toh [2010]. Similarly to the sparse logistic regression experiments, we solve each problem twice with different precision $\epsilon = 10^{-2}$ and $\epsilon = 10^{-6}$. Note that in all experiments QUIC consistently requires more flops to solve the problem than LHAC does. In one case where low precision is used it takes QUIC nearly three times more flops; even in the case QUIC is best at, where high precision is demanded, it can take QUIC 1.5 times more flops to achieve the same precision as LHAC did.

5 Conclusions

We have presented a general algorithm LHAC for efficiently using second-order information in training large-scale ℓ_1 -regularized convex models. We tested the algorithm on two instances of sparse logistic regression and five instances of sparse inverse covariance selection, and found that the algorithm is faster (sometimes overwhelmingly) than other specialized solvers on both models. The efficiency gains are due to two factors: the exploitation of the special structure present in the low-rank Hessian approximation matrix,

Table 2: Total complexity comparisons. p stands for dimension. ϵ indicates the optimization tolerance on the objective subgradient.

DATA SET	p	ϵ	COMPLEXITY($\times 10^6$)	
			QUIC	LHAC
LYMPH	587	10^{-2}	47	13
		10^{-6}	105	45
ER	692	10^{-2}	218	78
		10^{-6}	353	217
ARABIDOPSIS	834	10^{-2}	559	329
		10^{-6}	1001	804
LEUKEMIA	1255	10^{-2}	2101	574
		10^{-6}	3028	1326
HEREDITARY	1869	10^{-2}	16558	8333
		10^{-6}	18519	16770

and the greedy active-set strategy, which correctly identifies the non-zero features of the optimal solution.

References

- G. Andrew and J. Gao. Scalable training of l1-regularized log-linear models. *ICML*, pages 33–40, 2007.
- A. Beck and M. Teboulle. A Fast Iterative Shrinkage-Thresholding Algorithm for Linear Inverse Problems. *SIAM Journal on Imaging Sciences*, 2(1):183–202, 2009.
- R. E. Bixby, J. W. Gregory, I. J. Lustig, R. E. Marsten, and D. F. Shanno. Very large-scale linear programming: A case study in combining interior point and simplex methods. *Operations Research*, 40(5):pp. 885–897, 1992.
- R. Byrd, G. Chin, J. Nocedal, and F. Oztoprak. A family of second-order methods for convex l1-regularized optimization. *tech. rep.*, 2012.
- D. Donoho. De-noising by soft-thresholding. *Information Theory, IEEE Transactions on*, 41(3):613–627, may 1995.
- J. Friedman, T. Hastie, and R. Tibshirani. Sparse inverse covariance estimation with the graphical lasso. *Biostatistics Oxford England*, 9(3):432–41, 2008.
- J. Friedman, T. Hastie, and R. Tibshirani. Regularization paths for generalized linear models via coordinate descent. *Journal of Statistical Software*, 33:1–22, 2010.
- A. Genkin, D. D. Lewis, and D. Madigan. Large-scale bayesian logistic regression for text categorization. *Technometrics*, 49(3):291–304, 2007.
- C.-J. Hsieh, M. Sustik, I. Dhillon, and P. Ravikumar. Sparse inverse covariance matrix estimation using quadratic approximation. *NIPS*, 2011.

- D. D. Lewis, Y. Yang, T. Rose, and F. Li. Rcv1: A new benchmark collection for text categorization research. *JMLR*, 5:361–397, 2004.
- L. Li and K.-C. Toh. An inexact interior point method for L1-regularized sparse covariance selection. *Mathematical Programming*, 2:291–315, 2010.
- C.-J. Lin and J. Mor. Newtons method for large bound-constrained optimization problems. *SIAM Journal on Optimization*, 9(4):1100, 1999.
- Y. Nesterov. Gradient methods for minimizing composite objective function. *CORE report*, 2007.
- J. Nocedal and S. J. Wright. *Numerical Optimization*. Springer Series in Operations Research. Springer, New York, NY, USA, 2nd edition, 2006.
- P. A. Olsen, F. Oztoprak, J. Nocedal, and S. J. Rennie. Newton-Like Methods for Sparse Inverse Covariance Estimation, 2012.
- K. Scheinberg. An efficient implementation of an active set method for SVMs. *JMLR*, 7: 2237–2257, 2006.
- K. Scheinberg and I. Rish. SINCO - a greedy coordinate ascent method for sparse inverse covariance selection problem. *tech. rep.*, 2009.
- K. Scheinberg, S. Ma, and D. Goldfarb. Sparse inverse covariance selection via alternating linearization methods. *NIPS*, 2010.
- S. Shalev-Shwartz and A. Tewari. Stochastic methods for l1 regularized loss minimization. *ICML*, pages 929–936, 2009.
- R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society Series B Methodological*, 58(1):267–288, 1996.
- P. Tseng and S. Yun. A coordinate gradient descent method for nonsmooth separable minimization. *Mathematical Programming*, 117(1-2):387–423, Aug. 2009.
- S. Wright. Accelerated block-coordinate relaxation for regularized optimization. *SIAM Journal on Optimization*, 22(1):159–186, 2012.
- S. J. Wright, R. D. Nowak, and M. A. T. Figueiredo. Sparse reconstruction by separable approximation. *Trans. Sig. Proc.*, 57(7):2479–2493, July 2009.
- G.-X. Yuan, K.-W. Chang, C.-J. Hsieh, and C.-J. Lin. A comparison of optimization methods and software for large-scale l1-regularized linear classification. *JMLR*, 11: 3183–3234, 2010.
- G.-X. Yuan, C.-H. Ho, and C.-J. Lin. An improved GLMNET for l1-regularized logistic regression and support vector machines. *National Taiwan University*, 2011.